Ceph OSD node

- Ceph 3 (Luminous v12.x)

- 　　　　　　　　　　　　　　　　　　　.
- 　　　　　　　　　　　　　　　　　.

```
[root@monitor ~]# ceph df
[root@monitor ~]# rados df
[root@monitor ~]# ceph osd df
```

　　　　　　　　　　(I/O　　　　　　)

```
[root@monitor ~]# ceph osd set noscrub
[root@monitor ~]# ceph osd set nodeep-scrub
```

　　　　　　　　　(　　　OSD　　　　　　)

```
[root@monitor ~]# ceph daemon osd.0 config show

"osd_max_backfills": "1"
"osd_recovery_max_active": "3"
"osd_recovery_op_priority": "3"

[root@monitor ~]# ceph tell osd.* injectargs --osd-max-backfills 1 --osd-
recovery-max-active 1 --osd-recovery-op-priority 1
```

　　:
https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/3/html-single/configuration_guide/index#setting_a_specific_configuration_setting_at_runtime

OSD

OSD   OSD   active+clean   OSD
.

Ansible

OSD   ansible   shrink-osd.yml   shrink-osd-ceph-
disk.yml   .   osd_scenario   collocated   non-collocated   ,
shrink-osd-ceph-disk.yml   . osd_scenario   lvm   , shrink-osd.yml
          . admin keyring   .

1) ceph-ansible

```
[user@admin ~]$ cd /usr/share/ceph-ansible
```

2)   **yaml**

```
[root@admin ceph-ansible]# cp infrastructure-playbooks/shrink-osd.yml .

[root@admin ceph-ansible]# cp infrastructure-playbooks/shrink-osd-ceph-
disk.yml .
```

3)

```
ansible-playbook shrink-osd.yml -e osd_to_kill=$ID -u $ANSIBLE_USER
ansible-playbook shrink-osd-ceph-disk.yml -e osd_to_kill=$ID -u
$ANSIBLE_USER
```

- $ID :   OSD   .
- $ANSIBLE_USER : ansible

```
[user@admin ceph-ansible]$ ansible-playbook shrink-osd.yml -e osd_to_kill=1
-u user

[user@admin ceph-ansible]$ ansible-playbook shrink-osd-ceph-disk.yml -e
osd_to_kill=1 -u user
```

4) OSD

```
[root@mon ~]# ceph osd tree
```

CLI

root . . near
full

1) **OSD**

```
[root@osd ~]# ceph osd out 4
```

OSD active+clean OSD .

2) **OSD**

```
[root@osd ~]# systemctl disable ceph-osd@4
[root@osd ~]# systemctl stop ceph-osd@4
```

3) CRUSH map OSD

```
[root@osd ~]# ceph osd crush remove osd.4
```

4)

```
[root@osd ~]# ceph auth del osd.4
```

5) OSD

```
[root@osd ~]# ceph osd rm 4
```

6) **/etc/ceph.conf** **OSD**

```
[osd.4]
host = $HOST_NAME
```

7) /etc/fstab OSD

8) **/etc/ceph**

```
[root@osd ~]# scp /etc/ceph/ceph.conf root@node4:/etc/ceph/
```

OSD   .

```
[root@monitor ~]# ceph osd crush rm <        >
```

.

```
[root@monitor ~]# ceph daemon osd.0 config show
[root@monitor ~]# ceph tell osd.* injectargs --osd-max-backfills 1 --osd-
recovery-max-active 3 --osd-recovery-op-priority 3

[root@monitor ~]# ceph osd unset noscrub
[root@monitor ~]# ceph osd unset nodeep-scrub
```

- https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/3/html/operations_guide/handling-a-node-failure#removing-a-ceph-osd-node-ops